arXiv:1612.08228v1 [cs.DL] 25 Dec 2016

The misleading narrative of the canonical faculty productivity trajectory

Samuel F. Way,[1,*] Allison C. Morgan,[1,†] Aaron Clauset,[1,2,3,‡] and Daniel B. Larremore[3,§]

[1] *Department of Computer Science, University of Colorado, Boulder, CO, USA*
[2] *BioFrontiers Institute, University of Colorado, Boulder, CO, USA*
[3] *Santa Fe Institute, Santa Fe, NM, USA*

A researcher may publish tens or hundreds of papers, yet these contributions to the literature are not uniformly distributed over a career. Past analyses of the trajectories of faculty productivity suggest an intuitive and canonical pattern: after being hired, productivity tends to rise rapidly to an early peak and then gradually declines. Here, we test the universality of this conventional narrative by analyzing the structures of individual faculty productivity time series, constructed from over 200,000 publications matched with hiring data for 2453 tenure-track faculty in all 205 Ph.D-granting computer science departments in the U.S. and Canada. Unlike prior studies, which considered only some faculty or some institutions, or lacked common career reference points, here we combine a large bibliographic dataset with comprehensive information on career transitions that covers an entire field of study. We show that the conventional narrative describes only one third of faculty, regardless of department prestige, and the remaining two thirds of faculty exhibit a rich diversity of productivity patterns. To explain this diversity, we introduce a simple model of productivity trajectories, and explore which factors correlate with its parameters, showing that both individual productivity and the transition from first- to last-author publications correlate with departmental prestige.

INTRODUCTION

Scholarly publications serve as the primary mode of communication through which scientific knowledge is developed, discussed, and disseminated. The amount that an individual researcher contributes to this dialogue—their scholarly productivity—thus serves as an important measure of the rate at which they contribute units of knowledge to the field, and this measure is known to influence the placement of graduates into faculty jobs [1], the likelihood of being granted tenure [2, 3], and the ability to secure funding for future research [4].

The trajectory of productivity over the course of a researcher's lifetime has been studied for at least 60 years, with the common observation that a researcher's productivity rises rapidly to a peak and then slowly declines [5–9], inspiring the construction of mechanistic models with a similar profile [7, 9–11]. These models have included factors like cognitive decline with age, career age, finite supplies of human capital, knowledge advantages conferred by recent education, as well as skill deficits among the young, among others, and have been supported by the observation that productivity curves are not well described by even fourth-degree polynomial models [9]. Indeed, every study we found to date proposes or confirms a "rise and decline," "curvilinear," or "peak and tapering" productivity trajectory, regardless of whether researchers are binned by chronological age [5–8, 10–12], career age [9, 10], or (only for young researchers) years since first publication [13]. In fact, this conventional narrative of the life course is not restricted to academia, with similar trajectories observed in criminal behavior and artistic production in 1800s France [14] and even productivity of food acquisition by hunter-gatherers [15].

While these past studies have firmly established that the conventional academic productivity narrative is equally descriptive across fields and time, their analyses are based on averages over hundreds or thousands of individuals [5–15]. This raises two crucial and previously unanswered questions: is this average trajectory representative of individual faculty, and how much diversity is hidden by a focus on a central tendency over a population? To answer these questions, we combine and study two comprehensive datasets that span forty years of productivity for nearly every tenure-track professor in a North American Ph.D.-granting computer science department. By introducing a simple mathematical description, we map individuals' publication histories to a low-dimensional parameter space, revealing enormous diversity in the publication trends of individual faculty and showing that only a minority follow the conventional narrative of productivity. In fact, even among the conventional trajectories, individuals exhibit large fluctuations in their productivity around the average trend. Together, these results reveal that productivity patterns are both more diverse and less predictable than previously thought, and that population averages provide a dramatically inaccurate picture of intellectual contributions over time.

Moreover, while we show that the distribution of productivity trajectories resist natural categorizations, it is nevertheless possible to explore covariates that are associated with different regions of parameter space. The literature on such associations has avoided detailed trajectories and instead focused on the complicated relationship between prestige, productivity, and hiring. Past studies have found that researchers trained at prestigious

institutions are likely to remain productive [16, 17], regardless of where they place as faculty [18]. Other results link the prestige of the doctorate and the advisor to early-career productivity but not long-term productivity [19], at odds with others [20, 21] who found that early-career productivity predicts long-term productivity. Disagreement about hiring exists as well, with multiple studies finding that doctoral prestige, and not productivity drives the initial placement of faculty [22–24] while recent work based on comprehensive data in multiple fields suggests that prestige alone is insufficient to describe faculty placement [1, 25]. This, too, is complicated by hypotheses of mutually causality, where departments both select for and facilitate high productivity [26]. Unfortunately, while such studies shed light on a complicated system, they tend to restrict their analyses to unusual scientists, such as Nobel laureates or faculty at elite departments, rather than typical researchers. In contrast, the data analyzed here are comprehensive, covering faculty across the prestige hierarchy, and enable us to move beyond total productivity to study publication trajectories in light of prestige, hiring, and past productivity alike. To that end, we now describe the data used in our study.

DATA

This study exploits and combines two large datasets related to faculty productivity. The first is a comprehensive, hand-curated collection of education and academic appointment histories for tenure-track and tenured computer science faculty [25]. This dataset spans all 205 departmental or school-level academic units on the Computing Research Association's Forsythe List of Ph.D.-granting departments in computing-related disciplines in the United States and Canada [27]. For each department, the dataset provides a complete list of regular faculty for the 2011–2012 academic year, and for each of the 5032 faculty in this collection, it provides partial or complete information on their education and academic appointments, obtained from public online sources, mainly résumés and homepages. Of these, we selected the 2583 faculty who both received their Ph.D. from and held their first assistant professorship at one of these institutions, and for whom the year of that hire is known and occurred in 1970–2011. The first requirement ensured that we modeled the relatively closed North American faculty market; roughly 87% of computing faculty received their Ph.D. from one of the Forsythe institutions, and past analysis has shown that Canada and the United States are not distinct job markets in computer science [25]. A number of faculty were removed in this step because the location of their first assistant professorship was not known; these were mainly senior faculty.

The second dataset, constructed around the first, is a complete publication history as listed on DBLP, an online database [28] that provides open bibliographic information for most journals and conference proceedings relevant to computing research, using manual name disambiguation as necessary. For each paper in a faculty's publication history, we recorded the paper's title, author list (preserving author order), publication type (journal/conference paper or other), and year of publication. By following this procedure, we collected data for 200,476 publications which covered 2453 (95.0%) faculty in our sample.

Our combined dataset consists of the career trajectories of these 2453 tenure-track faculty as of 2011–2012: each professor's publicly-accessible metadata, their time-stamped Ph.D. and employment history, and the annotated time series of their publications. Briefly, we note that this dataset does not include information on faculty who have retired or left academia prior to 2011, and that publications are restricted to those in DBLP which focuses on computer science in particular.

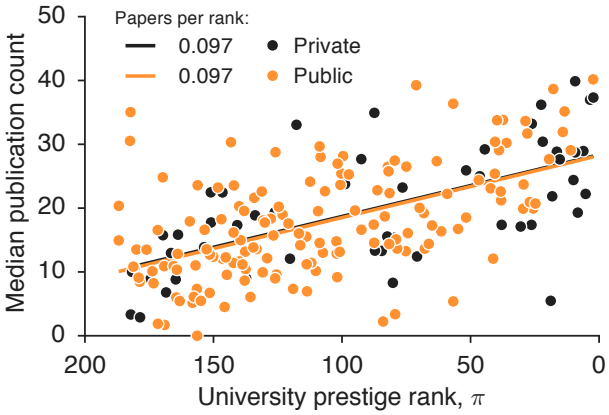

FIG. 1. **Publications correlate with prestige but not public/private status.** Dots indicate median number of publications over the early career (first five years) of researchers, adjusted for inflation (see text and Figure 2) for private (black) and public (orange) institutions, ordered by prestige [25]. Linear regressions and slopes are shown, indicating no difference between productivity at public and private institutions.

PRODUCTIVITY, PRESTIGE, AND TIME

Two broad trends characterize scholarly productivity in academic computer science. First, publication rates have been increasing over the past 45 years, and second, higher publication rates are correlated with higher prestige. These two observations are intertwined and underpin a number of subsequent analyses, so we explore them briefly in more depth.

Past studies have found that researchers at more prestigious institutions tend to be more productive [18, 19,

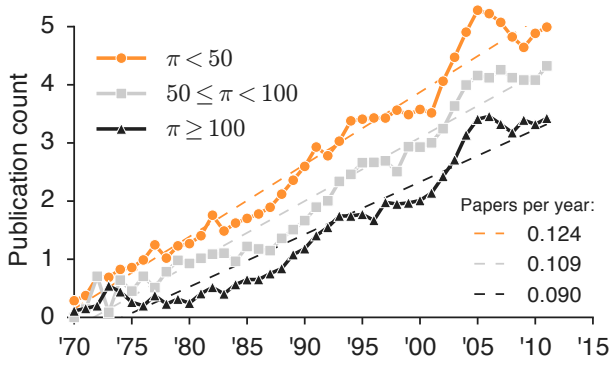

FIG. 2. **Annual publication rates have grown steadily.** For faculty in this study, per-person annual publications in the DBLP dataset have increase over time at a rate of approximately one additional paper every 9.6 years. This "rate of inflation" affects researchers at all levels of prestige rank π [25]. Slopes represent least-squares linear regressions.

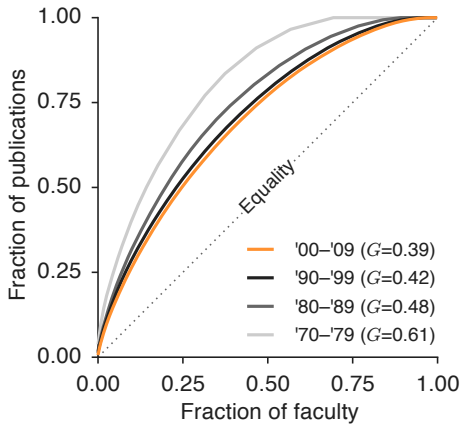

FIG. 3. **Production imbalance among faculty.** Lorenz curves of inflation-adjusted production by in-sample faculty, stratified by decade of first hire (see legend), show that approximately 20% of faculty account for half of all publications in the dataset. The Gini coefficient G for each curve is noted in the legend; diagonal line indicates equal production by all faculty.

23, 24, 26, 29]. Our data corroborate this finding, but we also find that the benefits of greater prestige hold regardless of whether an institution is public or private, for both early-career publications (first five years; Fig. 1) and lifetime publications. This result contradicts the conventional wisdom that private universities enjoy a productivity advantage over public ones. However, the conventional wisdom is likely skewed by a focus on elite departments, as 8 of the top 10 computer science departments are private [25], while in fact, private institutions are distributed evenly across all ranks. Regressing the median number of inflation-adjusted publications (see below) among faculty in a department against departmental prestige indicates that the relationships between prestige and productivity are statistically indistinguishable for public and private institutions, with expected increases in early career of roughly one publication for every 10-rank increase in prestige. A regression of productivity against prestige and a public/private dummy variable found a non-significant effect of public/private status ($p=0.878$, t-test). Expanding this analyis to include lifetime publications increases the prestige-publication slope to 3.5 publications per 10-rank increase in prestige but does not alter the non-significance of public/private status ($p=0.151$, t-test).

Past studies have also found that publication rates have increased over time [30, 31]. Knowing already that there are substantial differences in productivity by prestige, we separated universities arbitrarily into high, medium, and low prestige groups, and investigated whether the growth of publication rate varies by prestige. We find that the average number of publications per person produced in each calendar year has been increasing at all three strata of prestige at rates between 0.9 and 1.2 publications per decade, for 45 years (Figure 2). However, in addition to real increases in the number of papers published, some of this observed growth may be due to DBLP's relatively poorer coverage of older publications and venues, changes in what DBLP indexes relative to where researchers publish, and the growth of computer science as a field. Overall, the observed steady increase in productivity is not uniform across prestige, and the difference between production rates between higher and lower prestige departments has widened slightly but significantly over time ($p < 0.05$, t-test). In other words, prestigious and non-prestigious institutions have contributed to the overall "rate of inflation" at different rates. Not only are there small but significant differences in productivity by prestige (Figure 1) but those differences are slowly growing (Figure 2).

In order to investigate the productivity patterns of individual researchers, and test the conventional narrative of rapidly rising productivity followed by a gradual decline, for the remainder of this paper, we focus on time series of individual productivity. However, because of the observed growth in productivity, and its confounding with variations in DBLP coverage, it would be misleading to compare a 1975 publication with a 2005 publication. Thus, we will hereafter be using inflation-adjusted publications, correcting for the expansion of publication rates over time in the same way an economist would correct for currency inflation. All publications are reported in 2010-equivalent publications.

INDIVIDUALS

Examining the productivity trajectories of individual researchers, we find that they too exhibit substantial and

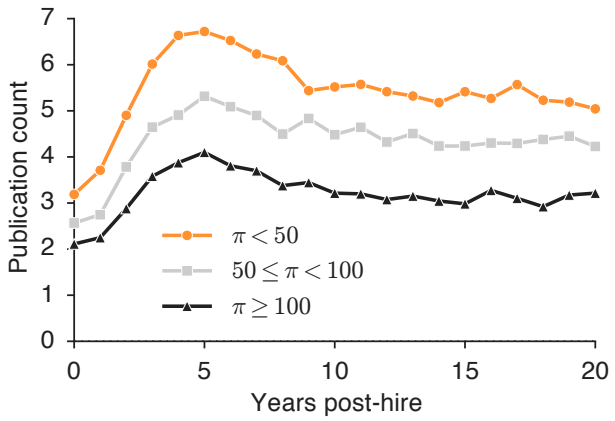

FIG. 4. **Average publications follow conventional narrative across prestige.** Three average productivity curves are shown, stratified by the prestige rank π of researchers' employing institutions as indicated. Averages over researchers at all levels of institutional prestige follow similar productivity trajectories, in agreement with the conventional narrative, but at differing scales of output.

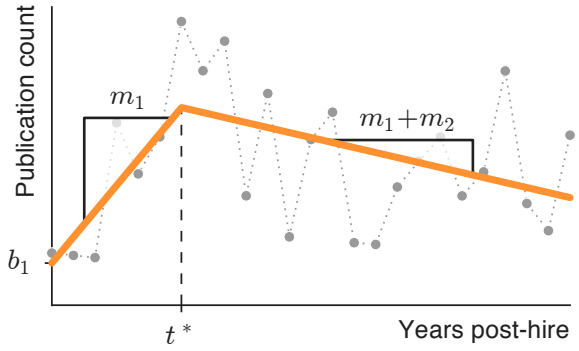

FIG. 5. **Example trajectory and piecewise model.** Dots represent empirical annual publications. Orange line shows best fit of piecewise linear model [Eq. (1)] with slopes m_1, $m_1 + m_2$, change point t^*, and intercept b_1 annotated.

significant differences in their publication rates. Early studies of scholarly productivity noted profound imbalance in the number of articles published by individual researchers [32, 33]. Cole [34] and Reskin [35] in the 1970s noted that about 50% of all scholarly articles were produced by about 15% of the scientific workforce. Our data reflect similar levels of imbalance, with approximately half of all contributions in the dataset authored by only 20% of all faculty. Figure 3 illustrates this inequality across researchers and time by showing the fraction Y of all publications in our sample that were produced by the most productive fraction X of all faculty (a Lorenz curve), for faculty first hired in each of the four decades that our data span. As indicated, the Gini coefficients for productivity imbalance have been declining, from 0.61 in the 1970s to 0.39 in the 2000s. This trend persists when researchers are restricted to only the publications within the first five years of their careers. There are several possible explanations for the trend of decreasing inequality in individual productivity. For instance, the lower end of the productivity distribution could have become relatively more productive over time, perhaps as more institutions shifted focus from teaching to research. Or, it may reflect a strengthening selective filter on highly productive faculty, perhaps as community expectations for continual productivity rose. It may also reflect non-uniform errors in the DBLP data, although our inflation correction should account for some of these. We leave the investigation of these possibilities for future work.

Instead, we focus on testing the conventional productivity narrative that has been described in various disciplines and at many points in time [5–15]: productivity climbs to a peak and then gradually declines over the course of the researcher's career. Across computer science faculty, we find that the average number of publications per year over a faculty career is highly stereotyped (Figure 4), with a rising productivity that peaks after around 5 years, declines slowly for another 5 years, and then remains roughly constant for any remaining years. Although departmental prestige correlates with productivity in several ways (Figures 1 and 2), it does not alter this stereotypical pattern, which appears essentially unchanged across departments with different levels of prestige, except for a roughly constant shift up as prestige increases (Figure 4).

The suggestion that productivity grows in the early years of a career has intuitive appeal. Professors settle into their research environments, begin training graduate students, and build their cases for promotion and tenure. Similarly, many reasons have been suggested for why productivity might decrease after promotion, including increased service and non-research commitments, declining cognitive abilities, and increased levels of distraction from outside work due to health issues and childcare obligations [36]. Although an average over faculty appears to confirm the stereotyped trajectory of rapid growth, peak, and slow decline, it does not reveal whether this average is representative of the many individual trajectories it averages over, nor does it show how much diversity there might be around the average, and whether that diversity correlates with other factors of interest.

To characterize the productivity pattern within an individual career, we fit a simple stereotypical model of productivity over time to the number of papers published per year,

$$f(t) = b_1 + m_1 t + m_2 (t - t^*) \cdot I(t > t^*) \ , \qquad (1)$$

a piecewise linear function in which t^* is the change point between the two lines, m_1 and $m_1 + m_2$ are the rates of

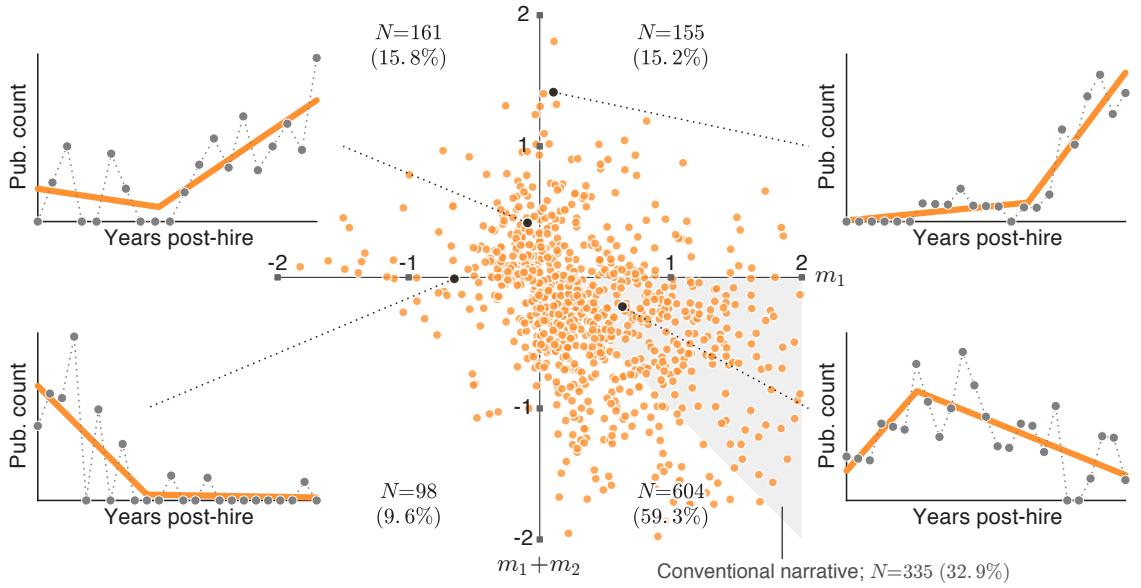

FIG. 6. **Distribution of individuals' productivity trajectory parameters.** The two piecewise slopes from the model in Eq. (1), m_1 and $m_1 + m_2$, for each individual in the dataset are distributed broadly across all four quadrants, indicated by counts and percentages. Subfigures show example publication trajectories to illustrate general characteristics of each quadrant. The shaded triangular region corresponds to the conventional narrative of early increase followed by gradual decline.

change in productivity before and after the change point, respectively, and b_1 is the initial productivity (Figure 5). By fitting these four parameters to each individual's publication trajectory, we map that individual's trajectory into a low-dimensional description of its overall pattern.

The narrative of "early growth in productivity, followed by a slow decline" implies four conditions on the inferred parameters: while the conditions of growth ($m_1 > 0$) and decline ($m_1+m_2 < 0$) are straightforward, we interpret "early growth" to mean that inferred peak productivity comes early ($t^* \leq 10$) and "slow" to mean that the slope of decline is smaller in magnitude that the slope of growth ($|m_1 + m_2| < m_1$). After fitting individual trajectory models to all 2543 faculty in our sample, we find that for faculty who have been employed for 10–25 years ($N = 1018$), only 29.7% follow the stereotypical trajectory. Even dropping the restriction on t^* increases the fraction meeting the stereotype to only 32.9%. In other words, the average trajectory, which has been held up as established fact for more than 50 years, describes the behavior of a minority of researchers, and instead most researchers follow qualitatively different trajectories.

Analyzing the signs of the two slope parameters, m_1 and m_1+m_2, publication trajectories are partitioned into one of four general classes, corresponding to the quadrants shown in Figure 6. We find enormous diversity in the shapes of individual trajectories, with nearly 10% falling into each major quadrant. Even among faculty whose publication rates grew and then declined (lower right quadrant, 59.3%), the conventional narrative only includes the 32.9% of individuals whose rate of growth exceeds their rate of decline ($m_1 > |m_1 + m_2|$, shaded region in Figure 6). To test the possibility that poor fits to very eccentric data might be driving these patterns, we restricted the data set to include the best 50% and 25% of model fits, but found no significant changes to the fractions of faculty in each quadrant.

The diverse cloud of faculty trajectory parameters shown in Figure 6 does not naturally separate into coherent clusters. In their absence, what are the covariates that predict which region of the plot an individual is likely to occupy? First, early-career growth rate of yearly publications m_1 is significantly higher for researchers at top-50 institutions ($p < 0.001$, Mann-Whitney), growing at 0.41 additional papers per year in the top 50 compared with 0.23 for others. Perhaps as a result—what goes up must come down—the slope after the point of inflection, m_1+m_2, is significantly more negative for researchers at higher-ranked institutions, compared to those at lower-ranked institutions ($p < 0.05$, Mann-Whitney). However, there was no significant difference in m_1 for researchers who received their doctorates at top-ranked institutions and those who trained at lower-ranked institutions, even when considering only those researchers employed at high prestige departments.

Second, contrasting with early-career growth rates, the early-career initial productivity b_1 is significantly higher both for those who are faculty in ($p < 0.001$, Mann-Whitney) and those who graduated from ($p <$

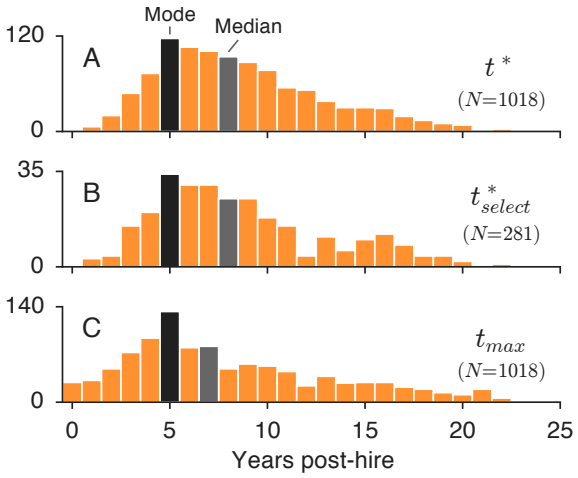

FIG. 7. **Productivity inflection points and maxima.** (A) Histogram of piecewise change points t^* for all individual faculty trajectories fitted with Eq. (1). (B) Histogram of t^* as in A, restricted to only those whose productivity trajectories are better modeled by Eq. (1) than a straight line, determined by AIC. (C) Histogram of maximum productivity year, measured directly from trajectories. Data reflect only those between 10 and 25 years of employment. Black bars indicate modes; grey bars indicate medians.

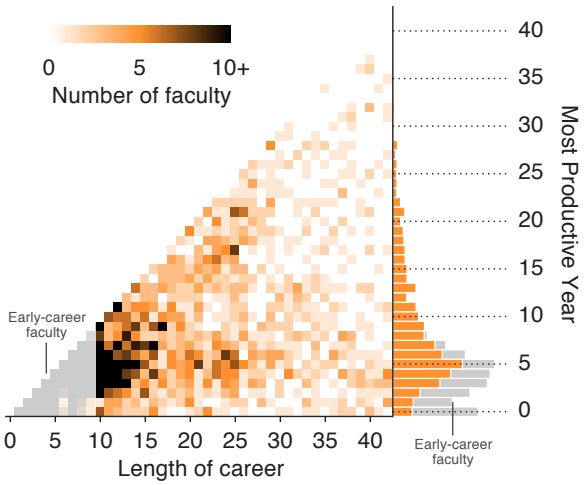

FIG. 8. **Heat map of each researchers' most productive years.** Each researcher's most productive year (empirically; not model fit) is plotted as a heatmap, sorted by the length of their career in our dataset. White box indicates researchers with fewer than 10 years of experience, whose most productive year is necessarily early. The marginal distribution (right) shows the empirically most productive year for all faculty in the dataset, separated by early career (first 10 years; grey) or later career (orange). The most common peak-productivity year is year 5, and only about half of senior faculty exhibit peak productivity in year 5 or earlier.

0.01, Mann-Whitney) top-50 departments. Unsurprisingly, we also find that researchers who have postdoctoral experience start out significantly more productive ($p < 0.05$, Mann-Whitney). These findings regarding m_1 and b_1 combine to suggest that current academic environment correlates with—and perhaps influences—productivity, while prior academic environment does not. Finally, faculty at top-ranked departments are statistically no more or less likely to be found within this triangular region, a result robust to alternatives cutoffs for "top ranked" institutions.

The inflection point within a career may indicate regime shifts in productivity, regardless of which type of trajectory an individual may follow. While the change-point parameter t^* does not correlate with the other parameters of $f(t)$, its distribution, shown in Figure 7A, reveals that for most faculty, the inferred change point in productivity rates occurs between years 5 and 10, with the median at 8 years. However, the piecewise $f(t)$ may be overparameterized for individuals with near-linear publication histories, running the risk of slightly overfitting their trajectories and making their fitted t^* values uninterpretable as productivity change points. To correct for this possibility, we performed model selection for all individuals, asking whether AIC with finite-size correction favored a straight line or the more complex $f(t)$. This process conservatively selected only 27.6% ($N = 281$) researchers who are more confidently modeled by the piecewise function, yet their distribution of productivity change points, shown in Figure 7B, is largely unaffected, with the modal value remaining at the 5th year, immediately preceding tenure decisions at most institutions. Nevertheless, there is still enormous diversity in career transitions, and the average remains misleading as the descriptor of a majority of individuals.

Unsurprisingly, a direct empirical examination of the set of publication time series reveals that a computer science professor's productivity is most likely to peak in the fifth year, shown in Figure 7C. However, a more detailed investigation shows that, in fact, peak productivity can occur in any year of a professor's career. Figure 8 translates each faculty member's career length (in our dataset) and most productive year relative to first assistant professorship into an ordered pair, creating a heat map of peak productivities. While the marginal distribution shows that 39.1% of faculty have their peak productivity within the first 6 years, with the modal peak year in year 5, there is substantial variance. Note, for example, that individuals along the bottom of the plot published the most in their first year as faculty, while individuals along the diagonal published the most in their last recorded year as faculty.

Finally, other transitions exist that are not quantifiable in publication counts alone, yet these are surprisingly well synchronized with the transitions noted above. As faculty train graduate students, their roles typically shift from lead researcher to senior advisor or principal

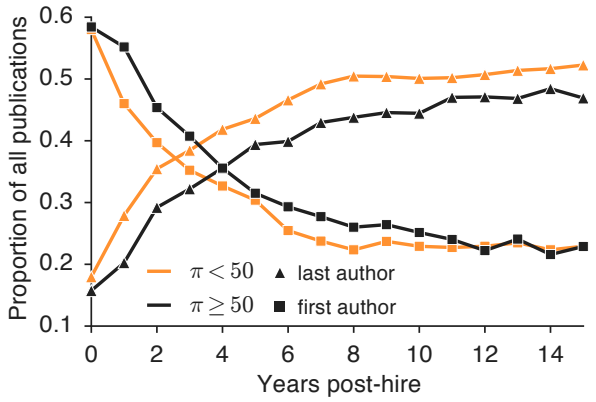

FIG. 9. **Early-career transitions in authorship roles.** The average proportion of first-author (squares) and last-author papers (triangles) as fraction of total, as a function of career age, split by top-50 institutions (orange) and others (black). On average, researchers at top-50 institutions transition more quickly into senior-authorship roles, compared to researchers at less prestigious institutions.

investigator, and this transition is typically reflected in a transition from first author to last author. Although this first/last convention is by no means universal—theoretical computer science, for example, orders authors alphabetically—we find evidence for a shift in roles in the data, with stability in the proportions of each type of publication around year eight, shown in Figure 9. Interestingly, the onset of this change is earlier among top-50 faculty, and their average proportion of last-author papers is significantly higher than those of other faculty, consistent with a hypothesis that faculty at top-50 institutions tend to begin working with students earlier, and have larger research groups.

DISCUSSION

The conventional narrative of faculty productivity over a career is pervasive, with repeated findings reinforcing a canonical trajectory where productivity rises rapidly to a peak early in one's career and then declines slowly [5–15]. This narrative shapes expectations of faculty across career stages, and unsurprisingly, publication counts have been shown to impact tenure decisions [2, 3], and the ability to secure funding for future research [4]. In this study, we showed that the conventional narrative, while intuitive, and certainly applicable to averages of many professors, is a remarkably inaccurate description of a typical professor's trajectory. By applying a simple piecewise-linear model to a comprehensive dataset of academic appointment histories and publication records, we found that only about one third of tenured or tenure-track computer science faculty resemble the average, irrespective of prestige.

While diverse, some aspects of a trajectory are nevertheless partially predictable. For example, although diversity of trajectories remains unaffected, productivity does tend to scale with prestige: researchers who graduated from or were hired by top-ranked institutions are significantly more productive at the onset of their careers, and, furthermore, productivity of high-prestige faculty tends to grow at faster rates and achieve higher peaks than researchers employed by other institutions. Together, these results support previously suggested hypotheses that top-ranked universities both select for and facilitate productivity [26]. In fact, our work hints at a possible causal relationship: researchers at top-ranked institutions appear to transition into leadership roles more quickly than others, further implicating facilitation effects in addition to selection.

Within the space of career trajectories, there is a noticeable tendency toward peak productivities and shifts in publication rates around 5–8 years. This is surely not a coincidence, given the fundamental role of tenure in an academic career. However, our data cannot yet say how, from a mechanistic perspective, the existence of tenure requirements drives faculty to change or shape their productivity before or after promotion. If anything, the results in this paper make clear that there are numerous ways in which computer scientists meet promotion requirements, not all of which necessarily involve publishing papers. Indeed, in parallel with career shapes more broadly, there remains enormous diversity in the distributions of productivity peaks and inflection points.

Computer science is, itself, a multifaceted field, and previous studies of the DBLP dataset revealed that productivity rates differ by subfield [1]. This observation, coupled with the zoo of fluctuating trajectories revealed in the current study, may suggest that year-to-year differences in individual trajectories are related to which subfields a researcher studies. Past work has revealed a first-mover advantage associated with entry into a rapidly growing field [37], so changes to individual research interests may contribute to noisy trajectories, particularly if they coincide with concentrated growth of popular new subfields.

The misleading narrative of the canonical productivity trajectory is not likely to be unique to computer science. Although future studies will reveal whether this claim is correct, we may need to rethink the conventional narrative of careers across academia. Other studies which investigate the impact of this pervasive narrative on decisions of promotion, retention, and funding would be particularly valuable. Expectations, whether perceived or enforced through tenure decisions, might give rise to some of our results. If these expectations vary from field to field, it is possible that while diversity remains a feature that spans academia, some types of trajectories may be more common in certain fields. Regardless of

whether future work combining productivity and hiring data across fields reveals this to be true, models of faculty productivity will need to be revisited and revised.

ACKNOWLEDGEMENTS

The authors thank Mirta Galesic for helpful conversations. All authors were supported by NSF award SMA 1633747; DBL was also supported by the Santa Fe Institute Omidyar Fellowship.

* samuel.way@colorado.edu
† allison.morgan@colorado.edu
‡ aaron.clauset@colorado.edu
§ larremore@santafe.edu

[1] S. F. Way, D. B. Larremore, and A. Clauset, in *Proceedings of the 25th International Conference on World Wide Web* (International World Wide Web Conferences Steering Committee, 2016) pp. 1169–1179.
[2] S. Cole and J. R. Cole, American Sociological Review, **32**, 377 (1967).
[3] W. F. Massy and A. K. Wilger, Change, **27**, 10 (1995), ISSN 00091383, 19399146.
[4] P. E. Stephan, Journal of Economic Literature, **34**, 1199 (1996).
[5] H. C. Lehman, The Scientific Monthly, **78**, 321 (1954).
[6] W. Dennis, Science, **123**, 724 (1956).
[7] A. M. Diamond, Journal of Gerontology, **41**, 520 (1986).
[8] K. L. Horner, J. P. Rushton, and P. A. Vernon, Psychology and Aging, **1**, 319 (1986).
[9] A. E. Bayer and J. E. Dutton, The Journal of Higher Education, **48**, 259 (1977).
[10] D. K. Simonton, Psychological Review, **104**, 66 (1997).
[11] S. G. Levin and P. E. Stephan, The American Economic Review, **81**, 114 (1991).
[12] S. Cole, American journal of sociology, 958 (1979).
[13] M. R. Symonds, N. J. Gemmell, T. L. Braisher, K. L. Gorringe, and M. A. Elgar, PLoS One, **1**, e127 (2006).
[14] A. Quetelet, *Sur l'homme et le développement de ses facultés ou essai de physique sociale*, Vol. 1 (Bachelier, 1835).
[15] H. Kaplan, K. Hill, J. Lancaster, and A. M. Hurtado, Evolutionary Anthropology Issues News and Reviews, **9**, 156 (2000).
[16] T. Caplow and R. J. McGee, *The Academic Marketplace* (Transaction Publishers, 1958).
[17] W. O. Hagstrom, "Departmental prestige and scientific productivity," Paper presented at the Annual Meeting of the American Sociological Association (1968).
[18] D. Crane, American Sociological Review, **30**, 699 (1965).
[19] B. F. Reskin, Sociology of Education, **52**, 129 (1979).
[20] J. S. Long, P. D. Allison, and R. McGinnis, American Sociological Review, **44**, 816 (1979).
[21] D. E. Chubin, A. L. Porter, and M. E. Boeckmann, American Sociological Review, **46**, 488 (1981).
[22] W. O. Hagstrom and L. L. Hargens, in *Cornell Conference on Human Mobility, Ithaca, NY* (1968).
[23] H. A. Zuckerman, Sociological Inquiry, **40**, 235 (1970).
[24] J. S. Long, American Sociological Review, **43**, 889 (1978).
[25] A. Clauset, S. Arbesman, and D. B. Larremore, Science Advances, **1**, e1400005 (2015).
[26] P. D. Allison and J. S. Long, American Sociological Review, **55**, 469 (1990).
[27] See http://archive.cra.org/reports/forsythe.html.
[28] See http://dblp.uni-trier.de.
[29] D. C. Pelz and F. M. Andrews, Ann Arbor, Mi: Institute for Social Research (1976).
[30] E. L. Dey, J. F. Milem, and J. B. Berger, Sociology of Education, **70**, 308 (1997).
[31] P. O. Larsen and M. Von Ins, Scientometrics, **84**, 575 (2010).
[32] A. J. Lotka, Journal of Washington Academy Sciences (1926).
[33] W. Shockley, Proceedings of the IRE, **45**, 279 (1957).
[34] J. R. Cole, *Fair science: Women in the Scientific Community* (New York: The Free Press, 1979).
[35] B. F. Reskin, American Journal of Sociology, **83**, 1235 (1978).
[36] M. F. Fox, Social Studies of Science, **13**, 285 (1983).
[37] M. E. Newman, EPL (Europhysics Letters), **86**, 68001 (2009).